Bargaining Solutions in a Social Network

Tanmoy Chakraborty and Michael Kearns

Department of Computer and Information Science
University of Pennsylvania

Abstract. We study the concept of bargaining solutions, which has been studied extensively in two-party settings, in a generalized setting involving arbitrary number of players and bilateral trade agreements over a social network. We define bargaining solutions in this setting, and show the existence of such solutions on all networks under some natural assumptions on the utility functions of the players. We also investigate the influence of network structure on equilibrium in our model, and note that approximate solutions can be computed efficiently when the networks are trees of bounded degree and the parties have nice utility functions.

1 Introduction

Bargaining has been studied extensively by economists and sociologists, and the most studied setup consists of two parties A and B, with utility functions $\mathcal{U}_A$ and $\mathcal{U}_B$, negotiating a bilateral deal. The deal, if agreed to by both parties, yield some fixed profit c. Such a scenario arises if two persons want to go into some business as partners. A and B also have *alternate options* α_A and α_B respectively, which is the amount of money they receive if the deal is not agreed upon. The negotiation involves how the profit from this deal is divided between the two parties. The final share that both parties agree to receive from the deal constitutes a *bargaining solution*.

Several bargaining solutions, which are predictions of how the profit will be shared, have been proposed by economists, the most well-known being the *Nash Bargaining Solution (NBS)* [8], which states that the bargaining solution finally adopted will be one that maximizes the product of the differential utilities of this deal to each party. The *differential utility* of A from the deal is the utility A receives by agreeing to the deal in excess of what it would receive without agreeing to the deal, that is, $\mathcal{U}_A(x) - \mathcal{U}_A(\alpha_A)$, where x is the share of profit A gets from the deal. Similarly, the differential utility of B is $\mathcal{U}_B(c-x) - \mathcal{U}_B(\alpha_b)$. NBS seeks to maximize $(\mathcal{U}_A(x) - \mathcal{U}_A(\alpha_A))(\mathcal{U}_B(c-x) - \mathcal{U}_B(\alpha_b))$.

Another extensively studied bargaining solution concept, known as the *Proportional Bargaining Solution (PBS)*, seeks to maximize the minimum of the differential utilities of the parties, that is, $\min\{\mathcal{U}_A(x) - \mathcal{U}_A(\alpha_A), \mathcal{U}_B(c-x) - \mathcal{U}_B(\alpha_b)\}$. There is a crucial axiomatic difference between the concepts of NBS and PBS – in fact, they are representatives from two broad classes of bargaining solution concepts that have been formulated and studied in literature (see Chapter 2 of

[1] for a discussion). One of the major axioms satisfied by NBS is that the bargaining solution should not be altered if the scale of the utility functions of the parties are altered by arbitrary constant factors. In other words, NBS is based on the axiom that utilities of different parties cannot be compared. However, this axiom is highly debated and several solutions that neglect this axiom, and instead choose to make *interpersonal comparison of utility*, have been proposed, and PBS is one of the most extensively studied among these solutions.

In this paper, we consider a generalization of the above two-party setting to a setting that involves arbitrary number of parties, but where the deals are still bilateral, and the alternate options are all zero. The parties shall be represented as vertices of a social network, where the edges represent bilateral deals that shall be negotiated, and weights on the edges represent the total profits from each deal. Thus the input to the problem, which we call the *network bargaining problem*, is an undirected graph with weights on edges, and an efficiently computable utility function for every vertex. Different deals may have different profits, which are represented by weights on edges of the input graph. A solution to the network bargaining problem is a prediction of how the profits on each edge is divided. We will primarily be interested in studying the effects of network topology on the solution, and so we shall often restrict our study to the case where the edges have unit weight, and all vertices have the same utility function. Effects of network topology on solutions of various network exchange models have been studied, theoretically as well as through human subject experiments ([5, 4, 2, 7]). Our goal is to develop a bargaining solution concept for the network bargaining problem, that will have a strong intuitive justification.

Braun and Gautschi [2] studied the network bargaining problem, and proposed a solution. Their solution is a direct generalization of the *weighted Nash Bargaining Solution* for bilateral deals. They assign a numerical bargaining power to each vertex based solely on its degree and the degrees of its neighbors, and also assume linear utility functions, and then negotiate each edge independently according to the bargaining powers. Kleinberg and Tardos [7] studied a variant of the network bargaining problem, which has the same input as our problem, but the solution has a restriction that each vertex can agree to a deal with at most one of its neighbors. They define an equilibrium-based solution, which they call a balanced outcome, where the agreement on every deal is required to meet a stability condition. The stability condition used for each edge in [7] is the NBS in a two-party setting with intuitively defined alternate options and linear utility functions. This model was also studied previously by Cook and Yamagishi [3]. In contrast to the model of [2], this model allows the equilibrium conditions and network topology to naturally exhibit bargaining power, instead of directly assigning a value.

Inspired by the notion of balanced outcomes in [3, 7], our solution for the network bargaining problem is also an equilibrium-based concept. We propose that the bargaining solution should be stable, and so no party should be keen on renegotiating a deal. For an edge $e = (u, v)$, we define the alternate options of u and v to be the total profits received by u and v, respectively, from other deals. The differential utility of the deal on e to u and v is now intuitively clear (see Section 2 for definition). Renegotiation occurs if the current division is not according to a standard two-party bargaining solution. In a bargaining state (specifying the division of profits on each edge), the deal on an edge is *stable* if the division of profits satisfy the two-party bargaining solution. We say that a state is an *equilibrium* if all edges are stable, and this is our bargaining solution. Depending on the two-party bargaining solution used for renegotiating edges, we have thus proposed two bargaining solutions, the *NBS equilibrium* and *PBS equilibrium*. These bargaining solutions are formally defined in Section 2.

In this paper, we completely characterize the PBS and NBS equilibria on every social network when all the vertices of the network have linear utility functions (functions with constant marginal utility). In this case, we show that there is a unique PBS equilibrium and a unique NBS equilibrium in every social network, and that the network topology has no influence on the solutions. A question crucial to the applicability of our model is to characterize structures in which there exists an equilibrium. We show that on any network, there exists a PBS equilibrium if all the utility functions are increasing and continuous. We also show that on any network, there exists an NBS equilibrium if all the utility functions are increasing, concave and twice differentiable.

The rest of this paper is organized as follows. Section 2 contains a formal introduction to the model and some basic lemmas that are applicable to two-party settings. Section 3 characterizes equilibria in our model when the utility functions are linear. Section 4 contains the proof of existence of PBS and NBS equilibria on all networks, for broad classes of utility functions. Section 5 provides some results about the effect of network structure on NBS equilibrium. And finally, Section 6 briefly describes an efficient algorithm to compute approximate PBS equilibria on trees with bounded degree and a specific utility function $\log(1 + x)$.

2 Preliminaries

The input to the *network bargaining problem* consists of an undirected graph $G(V, E)$ with n vertices and m weighted edges, where vertices represent people and edges represent possible bilateral trade deals, and a utility function $\mathcal{U}_v$ for each vertex v. The utility functions are all represented succintly and are computable in polynomial time.

Let $c(e)$ be the weight of an edge e in G. Let $e_1, e_2 \ldots e_m$ be an arbitrary ordering of the edges in $E(G)$, and let every edge be assigned an arbitrary direction, so that e_i is directed from endpoints u_i to v_i, $\forall i = 1, 2 \ldots m$. A *state of the bargaining model* is described by the division of profits on each edge of the graph. Let $x(u_i, e_i)$ and $x(v_i, e_i)$ denote the profits u_i and v_i receive from the agreement on the edge e_i, respectively. Note that $x(v_i, e_i) = c(e_i) - x(u_i, e_i)$. We shall represent a state of the bargaining model as a vector $s = (s_1, s_2 \ldots s_m) \in \mathbb{R}^m$ such that $s_i = x(u_i, e_i)$. Note that s uniquely determines the division of profits on all edges.

Definition 1. *Let $s \in \mathbb{R}^m$ be a state of the bargaining model for a graph G. For any vertex u and any edge e incident on u, let $\gamma_s(u)$ denote the total profit of a vertex u from all its deals with its neighbors. Let $x_s(u,e)$ denote the profit u gets from the agreement on edge e. Let $\alpha_s(u,e) = \gamma_s(u) - x_s(u,e)$ be the profit u receives from all its deals except that on e.*

If the current state of the bargaining model is s, and $e = (u,v)$ is renegotiated, then we say that $\alpha_s(u,e)$ and $\alpha_s(v,e)$ are the *alternate options* for u and v respectively, that is, the amount they receive if no agreement is reached on the deal on e. We shall drop the suffix s if we make a statement for any arbitrary state, or if the state is clear from the context.

Definition 2. *Let s be any state of the bargaining model. Let x be the profit of u from the deal on $e = (u,v)$. Then, the differential utility of u from this deal is $a_s(x) = \mathcal{U}_u(\alpha_s(u,e)+x) - \mathcal{U}_u(\alpha_s(u,e))$, and the differential utility of v from this deal is $b_s(x) = \mathcal{U}_v(\alpha_s(v,e)+c(e)-x) - \mathcal{U}_v(\alpha_s(v,e))$.*

Definition 3. *Let s be any state of the bargaining model. Define $y_s(u,e)$ to be the profit u would get on the edge $e = (u,v)$ if it is renegotiated (according to some two-party solution), the divisions on all other edges remaining unchanged. Also define update$(s,e) = |x_s(u,e) - y_s(u,e)|$.*

If e is renegotiated according to the *Nash Bargaining Solution* (NBS), then $y_s(u,e)$ is a value $0 \le x \le c(e)$ such that the NBS condition is satisfied, that is, the function $W_N(x) = a_s(x)b_s(x)$ is maximized. Instead, if e is renegotiated according to the *Proportional Bargaining Solution* (PBS), then $y_s(u,e)$ is a value $0 \le x \le c(e)$ such that the PBS condition is satisfied, that is, function $W_P(x) = \min\{a_s(x), b_s(x)\}$ is maximized.

The following lemmas give simpler equivalent conditions for PBS and NBS under certain assumptions about the utility functions, and are also applicable to the two-party setting.

Lemma 1. *If the utility functions of all vertices are increasing and continuous, then the PBS condition reduces to the condition $a_s(x) = b_s(x)$, and there is a unique solution x satisfying this condition.*

Proof. Since the utility functions are increasing, so $a_s(x)$ is an increasing function while $b_s(x)$ is a decreasing function, and both are non-negative functions when $x \in [0, c(e)]$. Also, $a_s(0) = b_s(c(e)) = 0$. Thus there is a unique value z such that $a_s(z) = b_s(z)$. If $x < z$, then $a_s(x) < b_s(x)$, and so $W_P(x) = a_s(x) < a_s(z)$, and if $x > z$, then $a_s(x) > b_s(x)$, and so $W_P(x) = b_s(x) < b_s(z)$. So $W_P(x)$ takes its maximum value only at $x = z$. □

A similar lemma holds for NBS when the utility functions are increasing, concave and twice differentiable, and is applicable to the two-party setting as well.

Lemma 2. *Let the utility functions of all vertices be increasing, concave and twice differentiable,. Moreover, let* $q_s(x) = \frac{a_s(x)}{a'_s(x)}$*, and let* $r_s(x) = -\frac{b_s(x)}{b'_s(x)}$*. Then the NBS condition reduces to* $q_s(x) = r_s(x)$*, and there is a unique solution* x *satisfying this condition.*

Proof. Differentiating $a_s(x)b_s(x)$ with respect to x and equating the derivative to zero, we get the equation

$$a_s(x)b'_s(x) + a'_s(x)b_s(x) = 0$$
$$\Rightarrow \frac{a'_s(x)}{a_s(x)} = -\frac{b'_s(x)}{b_s(x)}$$

Since the utility functions are concave and increasing, so is $a_s(x)$, while $b_s(x)$ is concave and decreasing, and both functions are positive. Thus we have $a_s(x) \geq 0,\ b_s(x) \geq 0,\ a'_s(x) > 0,\ b'_s(x) < 0, a''_s(x) < 0$, and $b''_s(x) < 0$. Differentiating $a_s(x)b_s(x)$ twice and using this information, we can easily see that this second derivative is always negative. Thus a zero of the first derivative is a maxima. Thus we have proved the equivalence of the two conditions.

To show uniqueness, let $q_s(x) = \frac{a_s(x)}{a'_s(x)}$, and let $r_s(x) = -\frac{b_s(x)}{b'_s(x)}$. Again, we can verify that $q'_s(x) > 0$ and $r'_s(x) < 0$ since $a''_s(x) < 0$ and $b''_s(x) < 0$. Moreover, $q_s(0) = 0$ and $r_s(c(e)) = 0$, since $a_s(0) = b_s(c(e)) = 0$. Thus, by an arugument similar to the proof of Lemma 1, we know that the increasing function $q_s(x)$ and the decreasing function $r_s(x)$ become equal at a unique value of x. □

Definition 4. *We say that an edge* e *is* stable *in a state* s *if renegotiating* e *does not change the division of profits on* e*, that is,* $update(s, e) = 0$*. We say that a state* s *is an equilibrium if all edges are stable. We say that* s *is an* ϵ*-approximate equilibrium if* $update(s, e) < \epsilon$ *for all edges* e*.*

We refer to an equilibrium as an *NBS equilibrium* if the renegotiations satisfy the NBS condition. We refer to the equilibrium as a *PBS equilibrium* if the renegotiations satisfy the PBS condition.

3 Linear Utility Functions: Characterizing All Equilibria

In this section, we characterize all possible NBS and PBS equilibria when all vertices have linear increasing utility functions, for every vertex v. Braun and Gautschi [2] make this assumption in their model, and so do Kleinberg and Tardos [7].

We show that in our model, if we make this assumption, there is a unique NBS equilibrium and a unique PBS equilibrium, and network topology has no influence on the division of profits on the deals at equilibrium. The following two theorems formalise these observations.

Theorem 1. *Suppose all vertices have linear increasing utility functions. Then there is a unique NBS equilibrium, in which the profit on every edge is divided equally between its two end-points.*

Proof. Let $\mathcal{U}_i(x) = k_i x + l_i \ \forall i \in V(G)$. Let s be an NBS equilibrium. Then, on any edge $e = (u, v)$, $a_s(x) = k_u x$ and $b_s(x) = k_v(c(e) - x)$. In particular, they are independent of $\alpha_s(u, e)$ and $\alpha_s(v, e)$. Since k_u and k_v are constants, the product $a_s(x)b_s(x)$ is maximized when $x = c(e)/2$. □

Theorem 2. *Suppose all vertices have linear increasing utility functions. Let $\mathcal{U}_i(x) = k_i x + l_i \ \forall i \in V(G)$. Then there is a unique PBS equilibrium, such that for any edge $e = (u, v)$, $x_s(u, e) = c(e)\frac{k_v}{k_u + k_v}$.*

Proof. Since the utility functions are increasing and continuous, we can apply Lemma 1 and for any edge $e = (u, v)$, solve the equation $a_s(x) = b_s(x)$. Substituting the values, we get the equation $k_u x = k_v(c(e) - x)$, which has the unique solution $x = c(e)\frac{k_v}{k_u + k_v}$. Again, the solution is independent of $\alpha_s(u, e)$ and $\alpha_s(v, e)$. Further, if $k_u = k_v$, then the solution becomes $x = c(e)/2$, which is the same as the solution for NBS. □

4 Existence of Equilibrium for General Utility Functions

We now turn our focus towards non-linear utility functions. In this section, we prove that PBS and NBS equilibria exist on all graphs, when the utility functions satisfy some natural conditions. The proofs use the Brouwer fixed point theorem, and is similar to the proof of existence of mixed Nash equilibrium in normal form games.

Theorem 3. *PBS equilibrium exists on any social network when all utility functions are increasing and continuous. NBS equilibrium exists on any social network when all utility functions are increasing, concave and twice differentiable.*

Essentially, it is sufficient for the utility functions to satisfy the following general condition of continuity:

Condition 1. *Let s be any state of the bargaining model, and $e = (u, v)$ be an edge. For every $\epsilon > 0$, there exists $\delta > 0$ such that for any state t such that $|\alpha_t(u, e) - \alpha_s(u, e)| < \delta$ and $|\alpha_t(v, e) - \alpha_s(v, e)| < \delta$, we have $|y_t(u, e) - y_s(u, e)| < \epsilon$.*

Note that $y_s(u, e)$ and $y_t(u, e)$ are influenced both by the utility functions as well as the two-party solution concept that is used (NBS or PBS). Thus whether Condition 1 holds will depend on whether the renegotiations follow the NBS or the PBS condition, and also on the utility functions.

Lemma 3. *If Condition 1 holds for the NBS solution concept or the PBS solution concept, then NBS or PBS equilibrium exists, respectively.*

Proof. We define a function $f : [0, 1]^m \to [0, 1]^m$ that maps every state s to another state $f(s)$. Given s, we can construct the unique solution t such that the deal on an edge $e = (u, v)$ in t is the renegotiated deal of e in s, that is,

$x_t(u,e) = y_s(u,e)$. We define $f(s)$ to be t. Thus, $f(s)$ is the "best-response" vector for s.

Clearly, s is an ϵ-approximate equilibrium if and only if $||s - f(s)||_\infty < \epsilon$. In particular, s is an equilibrium if and only if $f(s) = s$, that is, s is a fixed point of f. Also, $[0,1]^m$ is a closed, bounded and convex set. So if f were continuous, then we can immediately use Brouwer fixed point theorem to deduce that the equilibrium exists. Thus the following claim completes the proof. □

Claim. f is continuous if and only if Condition 1 holds.

Proof. Suppose Condition 1 holds for some ϵ and δ. Thus, if $||s - t||_\infty < \delta/n$, then $|\alpha_t(u,e) - \alpha_s(u,e)| < \delta$ and $|\alpha_t(v,e) - \alpha_s(v,e)| < \delta$, so, by Condition 1, $|y_t(u,e) - y_s(u,e)| < \epsilon$, and thus $||f(s) - f(t)||_\infty < \epsilon$. Since there exists a δ for every $\epsilon > 0$, so f is continuous.

Now suppose f is continuous. Let $\epsilon > 0$. Then there exists $\delta > 0$ such that for any solution t, if $||s - t||_\infty < \delta$, then $||f(s) - f(t)||_\infty < \epsilon$, which implies that coordinatewise for every edge e, we have $|y_t(u,e) - y_s(u,e)| < \epsilon$. Since this is true for all $\epsilon > 0$, Condition 1 holds. □

Lemma 4. *Condition 1 holds for all increasing, continuous utility functions when renegotiations follow the PBS condition.*

Proof. Let s be any state of the bargaining model and let $e = (u,v)$ be any edge. Here, Lemma 1 is applicable. Let $h(s,x) = a_s(x) - b_s(x)$. Also, let $g_s(x) = h(s,x)$ be a function defined on a particular state s. Note that g_s is an increasing, continuous function on the domain $[0, c(e)]$, $g_s(0) < 0$ and $g_s(c(e)) > 0$. The renegotiated value $y_s(u,e)$ is the unique zero of $g_s(x)$ between 0 and $c(e)$.

Let $y = y_s(u,e)$ be the zero of g_s. Let $\eta = \max\{|g_s(y-\epsilon)|, |g_s(y+\epsilon)|\}$. Then, since g_s is increasing, $\eta > 0$, and for all $x \in [0, c(e)] \setminus (y - \epsilon, y + \epsilon)$, $|g_s(x)| \geq \eta$.

Now, observe that $h(s,x)$ is dependent on $\alpha_s(u,e)$, $\alpha_s(v,e)$ and x only, and is continuous in all three of them when the utility functions are continuous. Thus, there exists $\delta > 0$ such that for any state t where $|\alpha_s(u,e) - \alpha_t(u,e)| < \delta$ and $|\alpha_s(v,e) - \alpha_t(v,e)| < \delta$, we have $|h(s,x) - h(t,x)| < \eta \ \forall x \in [0, c(e)]$, that is $|g_t(x) - g_s(x)| < \eta$. This implies that $g_t(x) \neq 0$ for all $x \in [0, c(e)] \setminus (y - \epsilon, y + \epsilon)$, and so the zero of g_t, which is $y_t(u,e)$, lies in the range $(y - \epsilon, y + \epsilon)$. □

Lemma 5. *Condition 1 holds for all increasing, concave and twice differentiable utility functions when renegotiations follow the NBS condition.*

Proof. Let s be any state of the bargaining model and let $e = (u,v)$ be any edge. Here, Lemma 2 is applicable. Let $h(s,x) = q_s(x) - r_s(x)$. Also, let $g_s(x) = h(s,x)$. The rest of the proof identically follows that of Lemma 4. □

Combining Lemmas 3, 4 and 5, we get Theorem 3.

5 Effect of Network Structure on NBS Equilibrium

In this section, we shall study the effect of network topology on NBS equilibrium. In the rest of this section, we shall assume here that all vertices have the *same utility function* $\mathcal{U}(x)$, and that the *deal on every edge has unit profit*, so that the network topology is solely responsible for any skewness in the distribution of profits in the NBS equilibrium. We also assume some natural properties of the utility function, and the following is our main result under these assumptions.

Theorem 4. *Let $\mathcal{U}(x)$ be the utility function of every vertex, and let all edges have unit weight. Let $\mathcal{U}(x)$ be increasing, twice differentiable and concave. Also, suppose $\frac{\mathcal{U}(x)-\mathcal{U}(0)}{\mathcal{U}'(x)} < Kx \ \forall x \in [0,1]$ for some constant K, and $|\mathcal{U}''(x)| \leq \epsilon(x)\mathcal{U}'(x)$ for some decreasing function $\epsilon(x)$. Let s be any NBS equilibrium in this network. Let $e=(u,v)$ be an edge such that u and v have degree more than $(K+1)d+1$ for some positive integer d. Then, $|x_s(u,e)-\frac{1}{2}| < \epsilon(d)$.*

Note that the assumptions on the utility function guarantee the existence of NBS equilibrium. Also note that the function $\mathcal{U}(x) = x^p$ for some $0 < p < 1$ satisfies the conditions of Theorem 4 with $K = p^{-1}$ and $\epsilon(x) = (1-p)/x$. The function $\mathcal{U}(x) = \log(1+x)$ satisfies the conditions of Theorem 4 as well, with $K = 2$, since $(1+x)\log(1+x) < (1+x)x \leq 2x$ when $x \in [0,1]$, and $\epsilon(x) = \frac{1}{1+x}$. To prove the above theorem, we will need the next two lemmas.

Lemma 6. *Let $\mathcal{U}(x)$ be increasing, twice differentiable and concave. Also, suppose $\frac{\mathcal{U}(x)-\mathcal{U}(0)}{\mathcal{U}'(x)} < Kx \ \forall x \in [0,1]$ Then at an NBS equilibrium s, for every edge $e=(u,v)$, $x_s(u,e) \geq \frac{1}{K+1}$ and $x_s(v,e) \geq \frac{1}{K+1}$.*

Proof. By Lemma 2, $x_s(u,e)$ is the unique solution to $q_s(x) = r_s(x)$. Note that $r_s(x) \geq 1-x$. This follows from the fact that $b_s(x) \geq -(1-x)b'_s(x)$, since $b_s(x)$ is concave and decreasing and $b_s(1) = 0$ (edges have unit weight).

Now, we show that $q_s(x) \leq Kx \ \forall x \in [0,1]$. To show this, we make the following claim.

Claim. If s_1 and s_2 are any two states such that $\alpha_{s_1}(u,e) > \alpha_{s_2}(u,e)$, then $q_{s_1}(x) < q_{s_2}(x) \ \forall x \in [0,1]$.

Proof (of Claim). To complete the proof, we need to prove our claim. For this, we view $q_s(x)$ as a function $z(\alpha)$ of $\alpha_s(u,e)$, keeping x constant, where $\alpha = \alpha_s(u,e)$, and observe that this function is decreasing. That is, $z(\alpha) = \frac{\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha)}{\mathcal{U}'(\alpha+x)}$, where x is constant. Note that $\mathcal{U}'(\alpha+x) - \mathcal{U}'(\alpha) < x\mathcal{U}''(\alpha+x)$, and $\mathcal{U}(\alpha+x) - \mathcal{U}(\alpha) > x\mathcal{U}'(\alpha+x)$. Differentiating $z(\alpha)$ with respect to α and using the above inequalities, we get that the numerator of the derivative $z'(\alpha)$ is $\mathcal{U}'(\alpha+x)(\mathcal{U}'(\alpha+x)-\mathcal{U}'(\alpha)) - (\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha))\mathcal{U}''(\alpha+x) < 0$, while the denominator is positive. So $z'(\alpha) < 0$, which implies our claim. □

Thus, $q_s(x)$ is greatest when $\alpha_s(u,e) = 0$, in which case $q_s(x) = \frac{\mathcal{U}(x)-\mathcal{U}(0)}{\mathcal{U}'(x)} < Kx$, the last inequality being our assumption. Thus, if x satisfies $q_s(x) = r_s(x)$, then $Kx > 1-x$, or $x > 1/(K+1)$. □

Lemma 7. *Let $\mathcal{U}(x)$ be increasing, twice differentiable and concave. Let s be an NBS equilibrium, $e = (u, v)$ be any edge, and $\epsilon > 0$. Also, let $|\mathcal{U}''(\alpha_s(u,e))| \leq \epsilon\mathcal{U}'(\alpha_s(u,e))$ and $|\mathcal{U}''(\alpha_s(v,e))| \leq \epsilon\mathcal{U}'(\alpha_s(v,e))$. Then, if u gets x on this agreement at equilibrium (and v gets $1-x$), then $|x - \frac{1}{2}| < \epsilon$.*

Proof. Let $\alpha = \alpha_s(u,e)$ and $\beta = \alpha_s(v,e)$. According to NBS, x is the solution maximizing $(\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha))(\mathcal{U}(\beta+1-x)-\mathcal{U}(\beta))$. Differentiating the quantity with respect to x and equating to zero, we get the following condition: $\mathcal{U}'(\alpha+x)(\mathcal{U}(\beta+1-x)-\mathcal{U}(\beta)) - \mathcal{U}'(\beta+1-x)(\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha)) = 0$.

We are seeking a solution to this equation. The left-hand-side of the above equation is a function $g(x)$ of the form $h_1(x) - h_2(x)$ that is decreasing, continuous, positive at $x = 0$ and negative at $x = 1$. Thus, the equation $g(x) = 0$ has a solution in $x \in (0,1)$.

Note that $\mathcal{U}(\beta+1-x)-\mathcal{U}(\beta)$ lies between $\mathcal{U}'(\beta)(1-x)$ and $\mathcal{U}'(\beta+1)(1-x)$, while $\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha)$ lies between $\mathcal{U}'(\alpha)x$ and $\mathcal{U}'(\alpha+1)x$. Also, $\mathcal{U}'(\beta+1-x) > \mathcal{U}'(\beta+1) > \mathcal{U}'(\beta)+\mathcal{U}''(\beta)$, and $\mathcal{U}'(\alpha+x) > \mathcal{U}'(\alpha+1) > \mathcal{U}'(\alpha)+\mathcal{U}''(\alpha)$, for all $x \in (0,1)$.

Now, for contradiction, suppose $x \leq \frac{1}{2} - \epsilon = \frac{1}{2}(1-2\epsilon)$. Then,

$$\begin{aligned}
h_1(x) &> \mathcal{U}'(\alpha+1)\mathcal{U}'(\beta+1)(1-x) > (\mathcal{U}'(\alpha)+\mathcal{U}''(\alpha))(\mathcal{U}'(\beta)+\mathcal{U}''(\beta))(1-x) \\
&= \mathcal{U}'(\alpha)\mathcal{U}'(\beta)(1+\frac{\mathcal{U}''(\alpha)}{\mathcal{U}'(\alpha)})(1+\frac{\mathcal{U}''(\beta)}{\mathcal{U}'(\beta)})(1-x) \\
&\geq \mathcal{U}'(\alpha)\mathcal{U}'(\beta)(1-\epsilon)(1-\epsilon)(\frac{1}{2})(1+2\epsilon) > \frac{1}{2}\mathcal{U}'(\alpha)\mathcal{U}'(\beta)(1-2\epsilon) \\
&\geq \mathcal{U}'(\alpha)\mathcal{U}'(\beta)x > h_2(x)
\end{aligned}$$

Thus, $g(x) > 0$ when $x \leq \frac{1}{2} - \epsilon$. Similarly, $g(x) < 0$ when $x \geq \frac{1}{2} + \epsilon$. □

Proof (of Theorem 4). There are $(K+1)d$ edges incident on each vertex u and v excluding (u,v), so Lemma 6 implies that at an NBS equilibrium, $\alpha_s(u,e) > \frac{1}{K+1}(K+1)d = d$ and $\alpha_s(v,e) > \frac{1}{K+1}(K+1)d = d$. Since $|\mathcal{U}''(x)| \leq \epsilon(x)\mathcal{U}'(x)$ and $\epsilon(x)$ is decreasing, we put $\epsilon = \epsilon(d) < \min\{\epsilon(\alpha_s(u,e)), \epsilon(\alpha_s(v,e))\}$ in Lemma 7 to obtain our result. □

6 Computing Approximate PBS Equilibria on Trees of Bounded Degree

In this section, as a first step towards settling the computational complexity of finding an equilibrium in our model, we note that approximate PBS equilibria can be computed efficiently when the networks are trees of bounded degree and utility function is same for all vertices and is very specific, as follows.

Theorem 5. *Suppose that the bargaining network is a tree with n vertices and maximum degree k, and weights on all edges bounded by C, and where all vertices have the same utility function $\mathcal{U}(x) = \log(1+x)$. There is an algorithm that computes an ϵ-approximate PBS equilibrium of this network in time $n(C\epsilon^{-1}k)^{O(k)}$.*

Since this algorithm is not central to this paper, and due to lack of space, we shall only provide its intuition and omit the details. Our algorithm is essentially a modification of the *TreeNash* algorithm of Kearns et. al. [6]. It is a dynamic programming technique on a rooted tree, where computation for the root u of a subtree can be easily completed if the same computation has been already completed for the children of u. The algorithm discretizes the division of profits on each edge to the multiples of some fraction $\delta = \epsilon/k$, and then computes a table for each subtree, under root u. A typical entry of the table stores whether there exists an approximate equilibrium in the subtree, given the total profit of u and its profit from the deal with its parent, and also the deals of u in at least one such equilibrium, if it exists.

Lemma 8 below is crucial for the correctness of our algorithm. It implies that the approximation factor achieved by the algorithm is proportional to the discretization factor δ. The lemma follows quite easily from Lemma 9. Lemma 9 depends heavily on the fact that the utility function is $\log(1 + x)$. However, similar results hold for many other utility functions, and our algorithm can be modified to apply to any such utility function.

Lemma 8. *Let s be an exact equilibrium on any graph of maximum degree at most k, and let $\mathcal{U}(x) = \log(1+x)$. Let t be any state with $l_\infty(s,t) = \max_{i=1}^m |s_i - s'_i| < \delta$. Then, t is a $k\delta$-approximate equilibrium.*

Lemma 9. *Let $\mathcal{U}(x) = \log(1+x)$. Let s and t be any two states, and $e = (u,v)$ be an edge, such that $|\alpha_t(u,e) - \alpha_s(u,e)| < \epsilon_1$ and $|\alpha_t(v,e) - \alpha_s(v,e)| < \epsilon_2$. If we use PBS for renegotiations, then $|y_t(u,e) - y_s(u,e)| < (\epsilon_1 + \epsilon_2)/2$.*

Proof. For any edge $e = (u,v)$, note that $|\alpha_t(u,e) - \alpha_s(u,e)| < (k-1)\delta$, and $|\alpha_t(u,e) - \alpha_s(u,e)| < (k-1)\delta$, since each edge other than e incident on u (or v, respectively) can contribute less than δ to either difference, and there are at most $k-1$ such edges. Thus, applying Lemma 9, we get that $|y_t(u,e) - y_s(u,e)| < (k-1)\delta$. Also, note that $x_s(u,e) = y_s(u,e)$, since s is an equilibrium. Thus

$$\begin{aligned} update(t,e) &= |y_t(u,e) - x_t(u,e)| \le |y_t(u,e) - y_s(u,e)| + |y_s(u,e) - x_t(u,e)| \\ &= |y_t(u,e) - y_s(u,e)| + |x_s(u,e) - x_t(u,e)| \\ &< (k-1)\delta + \delta = k\delta \end{aligned}$$

□

We shall now describe the algorithm. Given ϵ, let δ be the largest value such that $\delta < \epsilon/k$, and δ^{-1} is an integer. Our algorithm roots the tree at some vertex r, and maintains a two-dimensional table T_{uv} of order $(\delta^{-1}+1) \times (k\delta^{-1}+1)$ for every edge e between a parent u and its child v. For $i,j \in \mathbb{N},\ 0 \le i \le \delta^{-1},\ 0 \le j \le k\delta^{-1}$, $T_{uv}(i,j)$ is expected to hold the boolean answer to whether there exists an ϵ-approximate equilibrium *in the subtree under* v, where all profits are divided in multiples of δ^{-1}, given that v receives $i\delta$ profit from the edge $e = (u,v)$, such that v has a total profit of $j\delta$ from all its edges except e. Naturally, the subtree is said to be at equilibrium if all edges in it are stable. In addition, with each boolean value in the tables, there is also a pointer that are expected to point to a tuple of integers, indicating the various division of profits of v with its children, and the total profits of the children, that would give such an equilibrium. The pointer may point to multiple such tuples, each tuple giving such an equilibrium.

Initialize all values in all the tables to *false*. We fill up these tables with their correct values bottom-up, starting from the edges to the leaves. For edges $e = (u, v)$ such that v is a leaf, v's total profit is whatever it gets from e, so set $T_{uv}(i, i) = true$, for all $1 \leq i \leq \frac{1}{\delta}$. This completes the computation for the base case.

Now, for any edge $e = (u, v)$, suppose v has children $w_1, w_2 \dots w_l$, where $l \leq k$, and the tables $T_{vw_1}, T_{vw_2} \dots T_{vw_l}$ have already been correctly computed. We describe here a process to correctly compute the row $T_{uv}(i, .)$ for any i.

Pick integers $i_1, i_2, \dots i_l$ between 0 and δ^{-1}, and integers $j_1, j_2, \dots j_l$ between 0 and $k\delta^{-1}$, and for every such choice of a tuple of $2l$ integers, check the following:

1. Check if $T_{vw_q}(i_q, j_q)$ is *true*, for all $1 \leq q \leq l$.
2. Suppose $i\delta$ is the profit of v on edge e. Let $e_q = (v, w_q)$. Let w_q receive a profit of $j_q\delta$ from the edge e_q, and let $\alpha(w_q, e_q) = j_q\delta$, $\forall 1 \leq q \leq l$. Check if the edges e_q are all at ϵ-approximate equilibrium. If so, let $j = \sum_{q=1}^{l}(\delta^{-1} - i_q)$, and set $T_{uv}(i, j)$ to be true, store the tuple of $2l$ integers, and keep a pointer from $T_{uv}(i, j)$ to the tuple.

Once the tables for all the edges have been computed, we create an imaginary parent vertex a to the root r, and then compute the table T_{ar}. We are only interested in the entries of the column $T_{ar}(0, .)$, that is, equilibria where r receives zero profit from the imaginary edge (a, r).

Note that whenever a boolean value is true, the entry should also point to a tuple of integers giving the division of profits with, and total profits of its children. To extract an approximate equilibrium from these tables, we start from r and walk down the tree. We pick any j such that $T_{ar}(0, j)$ is true, and pick any tuple that it points to, say $i_1, i_2, \dots i_l, j_1, j_2 \dots j_l$ and if r has children $w_1, w_2 \dots w_l$, we construct part of the equilibrium by stating that r has a total profit of $j\delta$, it allows w_q to take a profit of i_q on the edge (r, w_q), and w_q has a total profit of $j_q\delta$ from its children. Then, we recursively look at the table values $T_{rw_1}(i_1, j_1), T_{rw_2}(i_2, j_2) \dots T_{rw_l}(i_l, j_l)$, which must all have the boolean value *true* (by the way the tables are filled up), and the tuples pointed to by these entries.

Correctness of the algorithm It is easy to see that the algorithm computes all ϵ-approximate equilibrium where the division of profits on all the edges are in multiples of δ^{-1}. By Lemma 8, since $k\delta < \epsilon$, we obtain that for every exact equilibrium, the algorithm will find at least one ϵ-approximate equilibrium that is less than $\delta < \epsilon$ distance away (under l_∞ norm) from the exact equilibrium. We complete the proof of Theorem 5 by analysing the running time of our algorithm.

Running Time of the algorithm While computing the table T_{uv}, on each guess of $2l$ integers, the algorithm spends $O(l)$ time. There are $O(\delta^{-1})$ choices for i_t, and

$O(k\delta^{-1})$ choices for j_t, and so there are at most $O((k\delta^{-2})^l)$ choices for the $2l$ integers. The time spent on computing the table T_{vw} is thus $O(l(k\delta^{-2})^l)$. Since $l < k$, and since δ is about ϵ/k, so the time to compute T_{vw} is at most $(k\epsilon^{-1})^{O(k)}$. Since there are $O(n)$ tables, the running time of our algorithm is bounded by $n(\epsilon^{-1}k)^{O(k)}$.

Finally, we would like to note that if all edges were not of unit weight, but all the weights were at most C, then the algorithm can be modified to run in time $n(C\epsilon^{-1}k)^{O(k)}$.

References

[1] K. Binmore. *Game Theory and the Social Contract, Volume 2: Just Playing*. The MIT Press, 1998.

[2] N. Braun and T. Gautschi. A nash bargaining model for simple exchange networks. *Social Networks*, 28(1):1–23, 2006.

[3] K. S. Cook and T. Yamagishi. Power in exchange networks: A power-dependence formulation. *Social Networks*, 14:245–265, 1992.

[4] S. Judd and M. Kearns. Behavioral experiments in networked trade. In *ACM Conference on Electronic Commerce*, 2008.

[5] S. M. Kakade, M. J. Kearns, L. E. Ortiz, R. Pemantle, and S. Suri. Economic properties of social networks. In *NIPS*, 2004.

[6] M. J. Kearns, M. L. Littman, and S. P. Singh. Graphical models for game theory. In *UAI*, pages 253–260, 2001.

[7] J. Kleinberg and E. Tardos. Balanced outcomes in social exchange networks. In *STOC*, 2008.

[8] J. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.